Spectral Analysis of Pollard Rho Collisions

Stephen D. Miller[*1] and Ramarathnam Venkatesan[2]

[1] Einstein Institute of Mathematics
The Hebrew University
Givat Ram, Jerusalem 91904, Israel
and
Department of Mathematics
Rutgers University
Piscataway, NJ 08854, USA
miller@math.huji.ac.il

[2] Microsoft Research
Cryptography and Anti-piracy Group
1 Microsoft Way, Redmond, WA 98052, USA
and
Cryptography Research Group
Microsoft Research India
Scientia - 196/36 2nd Main
Sadashivnagar, Bangalore 560 080, India
venkie@microsoft.com

Abstract. We show that the classical Pollard ρ algorithm for discrete logarithms produces a collision in expected time $O(\sqrt{n}(\log n)^3)$. This is the first nontrivial rigorous estimate for the collision probability for the *unaltered* Pollard ρ graph, and is close to the conjectured optimal bound of $O(\sqrt{n})$. The result is derived by showing that the mixing time for the random walk on this graph is $O((\log n)^3)$; without the squaring step in the Pollard ρ algorithm, the mixing time would be exponential in $\log n$. The technique involves a spectral analysis of *directed* graphs, which captures the effect of the squaring step.

Keywords: Pollard Rho algorithm, discrete logarithm, random walk, expander graph, collision time, mixing time, spectral analysis.

1 Introduction

Given a finite cyclic group G of order n and a generator g, the Discrete Logarithm Problem (DLOG) asks to invert the map $y \mapsto$

[*] Partially supported by NSF grant DMS-0301172 and an Alfred P. Sloan Foundation Fellowship.

g^y from $\mathbb{Z}/n\mathbb{Z}$ to G. Its presumed difficulty serves as the basis for several cryptosystems, most notably the Diffie-Hellman key exchange and some elliptic curve cryptosystems. Up to constant factors, the Pollard ρ algorithm is the most efficient and the only version with small memory known for solving DLOG on a general cyclic group – in particular for the group of points of an elliptic curve over a finite field.

We quickly recall the algorithm now. First one randomly partitions G into three sets S_1, S_2, and S_3. Set $x_0 = h$, or more generally to a random power $g^{r_1}h^{r_2}$. Given x_k, let $x_{k+1} = f(x_k)$, where $f : G \to G$ is defined by

$$f(x) \quad = \quad \begin{cases} gx\,, & x \in S_1\,; \\ hx\,, & x \in S_2\,; \\ x^2\,, & x \in S_3\,. \end{cases} \tag{1.1}$$

Repeat until a collision of values of the $\{x_k\}$ is detected (this is done using Floyd's method of comparing x_k to x_{2k}, which has the advantage of requiring minimal storage). We call the underlying directed graph in the above algorithm (whose vertices are the elements of G, and whose edges connect each vertex x to gx, hx, and x^2) as the *Pollard ρ Graph.* At each stage x_k may be written as $g^{a_k y+b_k}$, where $h = g^y$. The equality of x_k and x_ℓ means $a_k y + b_k = a_\ell y + b_\ell$, and solving for y (if possible) recovers the DLOG of $h = g^y$.

The above algorithm *heuristically* mimics a random walk. Were that indeed the case, a collision would be found in time $O(\sqrt{n})$, where n is the order of the group G. (The actual constant is more subtle; indeed, Teske [13] has given evidence that the walk is somewhat worse than random.)

The main result of this paper is the first rigorous nontrivial upper bound on the collision time. It is slightly worse than the conjectured $O(\sqrt{n})$, in that its runtime is $\widetilde{O}(\sqrt{n})$, i.e. off from $O(\sqrt{n})$ by at most a polynomial factor in $\log n$. As is standard and without any loss of generality, we tacitly make the following

$$\textbf{assumption:} \ \text{ the order } \ |G| = n \ \text{ is prime.} \tag{1.2}$$

Theorem 1.1. *Fix $\varepsilon > 0$. Then the Pollard ρ algorithm for discrete logarithms on G finds a collision in time $O_\varepsilon(\sqrt{n}\,(\log n)^3)$ with probability at least $1-\varepsilon$, where the probability is taken over all partitions of G into three sets S_1, S_2, and S_3.*

In the black-box group model (i.e. one which does not exploit any special properties of the encoding of group elements), a theorem of Shoup [11] states that any DLOG algorithm needs $\Omega(\sqrt{n})$ steps. Hence, aside from the probabilistic nature of the above algorithm and the extra factor of $(\log n)^3$, the estimate of Theorem 1.1 is sharp.

It should be noted that finding a collision does not necessarily imply finding a solution to DLOG; one must also show the resulting linear equation is nondegenerate. Since $n = |G|$ is prime this is believed to happen with overwhelming probability, much more so than for the above task of finding a collision in $O(\sqrt{n})$ time. This was shown for a variant of the Pollard ρ algorithm in [6], but the method there does not apply to the original algorithm itself. Using more refined techniques we are able to analyze this question further; the results of these investigations will be reported upon elsewhere.

This paper is the first analysis of the unmodified Pollard ρ Graph, including the fact that it is *directed*. One can obtain the required rapid mixing result for directed graphs by (a) assuming that rapid mixing holds for the undirected version, and (b) adding self-loops to each vertex. However, one still needs to prove (a), which in our situation is no simpler. In addition, the loops and loss of direction cause short cycles, which lead to awkward complications in the context of studying collisions.

Technically, analyzing directed graphs from a spectral point of view has the well known difficulty that a spectral gap is not equivalent to rapid mixing. A natural generalization of the spectral gap is the operator norm gap of the adjacency matrix, which suffices for our purposes (see Section 2). For a recent survey of mixing times on directed graphs, see [9].

The Pollard ρ graph is very similar to the graphs introduced by the authors in [8]. These graphs, which are related to expander graphs, also connect group elements x to $f(x)$ via the operations given in (1.1) – in particular they combine the operations of multiplication and squaring. The key estimate, a spectral bound on the adjacency operator on this graph, is used to show its random walks are rapidly mixing. Though the Pollard ρ walk is only *pseudorandom* (i.e., x_{k+1} is determined completely from x_k by its membership in S_1, S_2, or S_3), we are solely interested here in proving that it has a collision. The notions of random walk and pseudorandom walk (with random assignments of vertices in the sets S_i) coincide until a collision occurs.

1.1 Earlier Works

Previous experimental and theoretical studies of the Pollard ρ algorithm and its generalizations all came to the (unproven) conclusion that it runs in $O(\sqrt{n})$ time; this is in fact the basis for estimating the relative bit-for-bit security of elliptic curve cryptosystems compared to others, e.g. RSA. For an analysis of DLOG algorithms we refer the reader to the survey by Teske [14], and for an analysis of random walks on abelian groups, to the one by Hildebrand [4]. For the related Pollard ρ algorithm for factoring integers, Bach [1] improved the trivial bound of $O(n)$ by logarithmic factors.

An important statistic of the involved graphs is the *mixing time* τ, which loosely speaking is the amount of time needed for the random walk to converge to the uniform distribution, when started at an arbitrary node.[1] The existing approaches to modeling Pollard ρ can be grouped into two categories:

1. *Birthday attack in a totally random model:* each step is viewed as a move to a random group element, i.e. a completely random walk. In particular one assumes that the underlying graph has mixing time $\tau = 1$ and that its degree equals the group size; in reality the actual Pollard ρ graph has degree only 3. The $O(\sqrt{n})$ collision time is immediate for random walks of this sort.
2. *Random walk in an augmented graph*: The Pollard ρ graph is modified by increasing the number of generators k, but removing the squaring step. One then models the above transitions as random walks on directed abelian Cayley graphs. To ensure the mixing time is $\tau = O(\log|G|)$, however, the graph degree must grow at least logarithmically in $|G|$. The importance of τ stems from the fact that, typically, one incurs a overhead of multiplicative factor of τ^{const} in the overall algorithm.

[1] There are many inequivalent notions of mixing time (see [7]). Mixing time is only mentioned for purposes of rough comparison between different graphs; whatever we need about it is proved directly. Similarly, the reader need not recall any facts about expander graphs, which are mentioned only for motivation.

Teske [13], based on Hildebrand's results [4] on random walks on the cyclic group $\mathbb{Z}/m\mathbb{Z}$ with respect to steps of the form $x \mapsto x + a_i$, $i \leq k$, shows that the mixing time of an algorithm of the second type is on the order of $n^{\frac{2}{k-1}}$; she gives supporting numerics of random behavior for k large. In particular, without the squaring step the Pollard ρ walk would have mixing time on the order of n^2, well beyond the expected $O(\sqrt{n})$ collision time. This operation is an intriguing and cryptographically[2] important aspect of the Pollard ρ algorithm, and makes it inherently *non-abelian*: the Pollard ρ graphs are not isomorphic to any abelian Cayley graphs. Its effect cannot be accounted for by any analysis which studies only the additive structure of $\mathbb{Z}/m\mathbb{Z}$.

The present paper indeed analyzes the *exact* underlying Pollard ρ graph, without any modifications. We are able to show that the inclusion of the squaring step reduces the mixing time τ from exponential in $\log n$, to $O((\log n)^3)$ — see the remark following Proposition 3.2.

Our result and technique below easily generalize from the unmodified Pollard ρ algorithm, which has only 2 non-squaring operations, to the generalized algorithms proposed by Teske [13] which involve adding further such operations. Furthermore, it also applies more generally to additional powers other than squares. We omit the details, since the case of interest is in fact the most difficult, but have included a sketch of the argument at the end of the paper.

2 Rapid mixing on directed graphs

In the next two sections we will describe some results in graph theory which are needed for the proof of Theorem 1.1. Some of this material is analogous to known results for *undirected* graphs (see, for example, [2]); however, since the literature on spectral analytic aspects of directed graphs is relatively scarce, we have decided to give full proofs for completeness.

The three properties of subset expansion, spectral gap, and rapid mixing are all equivalent for families of undirected graphs with fixed

[2] In this version one can derive a secure hash function [5] whose security is based on the difficulty of the discrete logarithm problem; here the input describes the path taken in the graph from a fixed node, and the hash value is the end point.

degree. This equivalence, however, fails for directed graphs. Although a result of Fill [3] allows one to deduce rapid mixing on directed graphs from undirected analogs, it involves adding self-loops (which the Pollard ρ graph does not have) and some additional overhead. In any event, it requires proving an estimate about the spectrum of the undirected graph. We are able to use the inequality [8, (A.10)], which came up in studying related undirected graphs, in order to give a bound on the operator norm of the directed graphs. This bound, combined with Lemma 2.1, gives an estimate of $\tau = O((\log n)^3)$ for the mixing time of the Pollard ρ graph.

Let $\mathit{\Gamma}$ denote a graph with a finite set of vertices V and edges E. Our graphs will be directed graphs, meaning that each edge has an orientation; an edge from v_1 to v_2 will be denoted by $v_1 \to v_2$. Assume that $\mathit{\Gamma}$ has *degree* k, in other words that each vertex has exactly k edges coming in and k edges coming out of it. The *adjacency* operator A acts on $L^2(V) = \{f : V \to \mathbb{C}\}$ by summing over these k neighbors:

$$(Af)(v) \quad = \quad \sum_{v \to w} f(w)\,. \tag{2.1}$$

Clearly constant functions, such as $\mathbb{1}(v) \equiv 1$, are eigenfunctions of A with eigenvalue k. Accordingly, $\mathbb{1}$ is termed the trivial eigenfunction and k the trivial eigenvalue of A. Representing A as a $|V| \times |V|$ matrix, we see it has exactly k ones in each row and column, with all other entries equal to zero. It follows that $\mathbb{1}$ is also an eigenfunction with eigenvalue k of the adjoint operator A^*

$$(A^* f)(v) \quad = \quad \sum_{w \to v} f(w)\,, \tag{2.2}$$

and that all eigenvalues λ of A or A^* satisfy the bound $|\lambda| \le k$.

The subject of *expander graphs* is concerned with bounding the (undirected) adjacency operator's restriction to the subspace $L_0 = \{f \in L^2(V) \mid f \perp \mathbb{1}\}$, i.e. the orthogonal complement of the constant functions under the L^2-inner product. This is customarily done by bounding the nontrivial eigenvalues away from k. However, since the adjacency operator A of a directed graph might not be self-adjoint, the operator norm can sometimes be a more useful quantity to study. We next state a lemma relating it to the rapid mixing of the random walk. To put the statement into perspective, consider the k^r random walks on Γ of length r starting from any fixed vertex. One expects a uniformly distributed walk to land in any fixed subset S with probability roughly $\frac{|S|}{|V|}$. The lemma gives a condition on the operator norm for this probability to in fact lie between $\frac{1}{2}\frac{|S|}{|V|}$ and $\frac{3}{2}\frac{|S|}{|V|}$ for moderately large values of r. This can alternatively be thought of as giving an upper bound on the mixing time.

Lemma 2.1. *Let Γ denote a directed graph of degree k on n vertices. Suppose that there exists a constant $\mu < k$ such that $\|Af\| \leq \mu \|f\|$ for all $f \in L^2(V)$ such that $f \perp \mathbb{1}$. Let S be an arbitrary subset of V. Then the number of paths of length $r \geq \frac{\log(2n)}{\log(k/\mu)}$ which start from any given vertex and end in S is between $\frac{1}{2}k^r\frac{|S|}{|V|}$ and $\frac{3}{2}k^r\frac{|S|}{|V|}$.*

Proof. Let y denote an arbitrary vertex in V, and χ_S and $\chi_{\{y\}}$ the characteristic functions of S and $\{y\}$, respectively. The number of paths of length r starting at y and ending in S is exactly the $L^2(V)$-inner product $\langle \chi_S, A^r \chi_{\{y\}} \rangle$. Write

$$\chi_S \quad = \quad \frac{|S|}{n}\mathbb{1} + w \quad \text{and} \quad \chi_{\{y\}} \quad = \quad \frac{1}{n}\mathbb{1} + u\,, \tag{2.3}$$

where $w, u \perp \mathbb{1}$. Because $\mathbb{1}$ is an eigenfunction of A^*, A preserves the orthogonal complement of $\mathbb{1}$, and thus

$$\|A^r u\| \quad \leq \quad \mu\,\|A^{r-1}u\| \quad \leq \quad \cdots \quad \leq \quad \mu^r\,\|u\|\,. \tag{2.4}$$

Also, by orthogonality

$$\|w\| \;\leq\; \|\chi_S\| \;=\; \sqrt{|S|} \qquad \text{and} \qquad \|u\| \;\leq\; \|\chi_{\{y\}}\| \;=\; 1\,. \tag{2.5}$$

We have that $A^r\chi_{\{y\}} = \frac{1}{n}k^r\mathbb{1} + A^r u$, so the inner product may be calculated as

$$\langle \chi_S, A^r\chi_{\{y\}} \rangle = \frac{|S|}{n}k^r + \langle w, A^r u \rangle\,. \tag{2.6}$$

It now suffices to show that the absolute value of the second term on the righthand side is bounded by half of the first term. Indeed,

$$|\langle w, A^r u \rangle| \quad \leq \quad \|w\|\,\|A^r u\| \quad \leq \quad \mu^r\sqrt{|S|}\,, \tag{2.7}$$

and

$$\mu^r \sqrt{|S|} \;\le\; \frac{1}{2n} k^r \sqrt{|S|} \;\le\; \frac{1}{2} k^r \frac{|S|}{n} \tag{2.8}$$

when $r \ge \frac{\log(2n)}{\log(k/\mu)}$. □

3 Collisions on the Pollard ρ graph

In this section, we prove an operator norm bound on the Pollard ρ graph that is later used in conjunction with Lemma 2.1. These graphs are closely related to an undirected graph studied in [8, Theorem 4.1]. We will start by quoting a special case of the key estimate of that paper, which concerns quadratic forms. At first glance, the analysis is reminiscent of the of the Hilbert inequality from analytic number theory (see [10, 12]), but where the quadratic form coefficients are expressed as $1/\sin(\mu_j - \mu_k)$.

Let n be an odd integer and $\lambda_k = |\cos(\pi k/n)|$ for $k \in \mathbb{Z}/n\mathbb{Z}$. Consider the quadratic form $Q : \mathbb{R}^{n-1} \to \mathbb{R}$ given by

$$Q(x_1, \ldots, x_{n-1}) \;\; := \;\; \sum_{k=1}^{n-1} x_k \, x_{2k} \, \lambda_k \, , \tag{3.1}$$

in which the subscripts are interpreted modulo n.

Proposition 3.1. *There exists an absolute constant $c > 0$ such that*

$$|Q(x_1, \ldots, x_{n-1})| \;\; \le \;\; \left(1 - \frac{c}{(\log n)^2}\right) \sum_{k=1}^{n-1} x_k^2 \, . \tag{3.2}$$

Proof. Let γ_k be arbitrary positive quantities (which will be specified later in the proof). Since

$$\gamma_k \, x_k^2 + \gamma_k^{-1} \, x_{2k}^2 \pm 2 \, x_k \, x_{2k} \;\; = \;\; \left(\gamma_k^{1/2} \, x_k \;\pm\; \gamma_k^{-1/2} \, x_{2k}\right)^2 \;\ge\; 0 \, , \tag{3.3}$$

one has that

$$|Q(\boldsymbol{x})| \;\le\; \frac{1}{2} \sum_{k=1}^{n-1} \left(\gamma_k \, x_k^2 \;+\; \gamma_k^{-1} \, x_{2k}^2\right) \, \lambda_k \;\; = \;\; \frac{1}{2} \sum_{k=1}^{n-1} x_k^2 \left(\gamma_k \, \lambda_k + \gamma_{\bar{2}k}^{-1} \lambda_{\bar{2}k}\right) , \tag{3.4}$$

where $\bar{2}$ denotes the multiplicative inverse to 2 modulo n. The proposition follows if we can choose γ_k and an absolute constant $c > 0$ such that

$$\gamma_k \, \lambda_k \ + \ \gamma_{\bar{2}k}^{-1} \, \lambda_{\bar{2}k} \ < \ 2 \ - \ \frac{c}{(\log n)^2} \qquad \text{for all } 1 \leq k < n \, . \tag{3.5}$$

Now we come to the definition of the γ_k. We set $\gamma_k = 1$ for $n/4 \leq k \leq 3n/4$; the definition for the set of other nonzero indices $\mathcal{S}$ is more involved. For $s \geq 0$, define

$$t_s \ = \ 1 \ - \ s \, \frac{d}{(\log n)^2} \, ,$$

where $d > 0$ is a small constant that is chosen at the end of the proof. Given an integer ℓ in the range $-n/4 < \ell < n/4$, we define $u(\ell)$ to be order to which 2 divides ℓ. For the residues $k \in \mathcal{S}$, which are all equivalent modulo n to some integer ℓ in the interval $-n/4 < \ell < n/4$, we define $\gamma_k = t_{u(\ell)}$. Note also that $\lambda_k \leq 1/\sqrt{2}$ for $k \notin \mathcal{S}$, and is always ≤ 1. With these choices the lefthand side of (3.5) is bounded by

$$\gamma_k \, \lambda_k \ + \ \gamma_{\bar{2}k}^{-1} \, \lambda_{\bar{2}k} \ \leq \ \begin{cases} \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{2}} \, , & k, \bar{2}k \notin \mathcal{S} \\ \frac{1}{\sqrt{2}} + \gamma_{\bar{2}k}^{-1} \, , & k \notin \mathcal{S}, \, \bar{2}k \in \mathcal{S} \\ \gamma_k + \frac{1}{\sqrt{2}} \, , & k \in \mathcal{S}, \, \bar{2}k \notin \mathcal{S} \\ \gamma_k + \gamma_{\bar{2}k}^{-1} \, , & k, \bar{2}k \in \mathcal{S}. \end{cases} \tag{3.6}$$

In the last case, the residues k and $\bar{2}k$ both lie in $\mathcal{S}$. The integer $\ell \equiv \bar{2}k \pmod{n}$, $-n/4 < \ell < n/4$, of course satisfies the congruence $2\ell \equiv k \pmod{n}$. Since $k \in \mathcal{S}$, 2ℓ is the unique integer in $(-n/4, n/4)$ congruent to k. That means $\gamma_k = t_{s+1}$ and $\gamma_{\bar{2}k} = t_s$ for some positive integer $s = O(\log n)$. A bound for the last case in (3.6) is therefore $t_{s+1} + t_s^{-1} = 2 - d/(\log n)^2 + O(s^2 d^2/(\log n)^4)$. We conclude in each of the four cases that, for d sufficiently small, there exists a positive constant $c > 0$ such that (3.5) holds. □

The Pollard ρ graph, introduced earlier, is the graph on $\mathbb{Z}/n\mathbb{Z}$ whose edges represent the possibilities involved in applying the iterating function (1.1):

$$\Gamma \text{ has vertices } V = \mathbb{Z}/n\mathbb{Z} \text{ and directed edges } x \to x+1,\ x \to x+y, \\ \text{and } x \to 2x \text{ for each } x \in V \text{ (where } y \neq 1) \ . \tag{3.7}$$

Proposition 3.2. *Let A denote the adjacency operator of the graph (3.7) and assume that n is prime. Then there exists an absolute constant $c > 0$ such that*

$$\|Af\| \quad \leq \quad \left(3 - \frac{c}{(\log n)^2}\right) \|f\| \tag{3.8}$$

for all $f \in L^2(V)$ such that $f \perp \mathbb{1}$.

Proof. Let $\chi_k : \mathbb{Z}/n\mathbb{Z} \to \mathbb{C}$ denote the additive character given by $\chi_k(x) = e^{2\pi i k x/n}$. These characters, for $1 \leq k < n$, form a basis of functions $L_0 = \{f \in L^2 \mid f \perp \mathbb{1}\}$. The action of A on this basis is given by

$$A\,\chi_k \quad = \quad d_k\,\chi_k + \chi_{2k} \quad , \qquad \text{where } \ d_k \ = \ e^{2\pi i k/n} + e^{2\pi i k y/n} \, . \tag{3.9}$$

One has that $|d_k| \ = \ 2|\cos(\frac{\pi k(y-1)}{n})| \ = \ 2\lambda_{k(y-1)}$. Using the inner product relation

$$\langle \chi_k, \chi_\ell \rangle \quad = \quad \begin{cases} n\,, & k = \ell \\ 0\ , & \text{otherwise}\,, \end{cases} \tag{3.10}$$

we compute that $\|f\|^2 = n \sum |c_k|^2$, where $f = \sum_{k \neq 0} c_k \chi_k$. Likewise,

$$\begin{aligned} \|Af\|^2 \quad &= \quad \langle Af, Af \rangle \quad = \\ \sum_{k,\ell \neq 0} c_k\,\overline{c_\ell}\, &[\langle d_k\chi_k, d_\ell\chi_\ell\rangle + \langle \chi_{2k}, \chi_{2\ell}\rangle + \langle d_k\chi_k, \chi_{2\ell}\rangle + \langle \chi_{2k}, d_\ell\chi_\ell\rangle] \\ &\leq \quad n\left(5\sum |c_k|^2 + 2\sum |c_k||c_{2k}||d_{2k}|\right) . \end{aligned} \tag{3.11}$$

Note that $|d_k| = 2\lambda_{k(y-1)}$, and that $y-1$ and 2 are invertible in $\mathbb{Z}/n\mathbb{Z}$, by assumption in (3.7). The result now follows from (3.2) with the choice of $x_{2(y-1)k} = |c_k|$.

□

Remark: the above Proposition, in combination with Lemma 2.1, is the source of the $\tau = O((\log n)^3)$ mixing time estimate for the Pollard ρ graph that we mentioned in the introduction.

Proof (of Theorem 1.1). Consider the set S of the first $t = \lfloor \sqrt{n} \rfloor$ iterates $x_1, x_2, \ldots, x_t$. We may assume that $|S| = t$, for otherwise a collision has already occurred in the first $\sqrt{n}$ steps. Lemma 2.1 and Proposition 3.2 show that the probability of a walk of length $r \gg (\log n)^3$ reaching S from any fixed vertex is at least $1/(2\sqrt{n})$. Thus the probabilities that $x_{t+r}, x_{t+2r}, x_{t+3r}, \ldots, x_{t+kr}$ lie in S are all, independently, at least $1/(3t)$. One concludes that for k on the order of $3bt$, b fixed, the probability that none of these points lies in S is at most $(1 - \frac{1}{3t})^{3bt} \approx e^{-b}$, which is less than ε for large values of b.

Generalizations: the analysis presented here extends to generalized Pollard ρ graphs in which each vertex x is connected to others of the form xg_i, for various group elements g_i, along with powers x^{r_j}. This can be done as follows. First of all, if r-th powers are to be used instead of squares, then the subscript $2k$ in (3.1) must be changed to rk. The key bound on (3.2), stated here for $r = 2$, in fact holds for any fixed integer $r > 1$ which is relatively prime to n [8, Appendix]. Thus changing the squaring step to $x \rightarrow x^r$ does not change the end results. Secondly, the proof of the bound (3.8) requires only some cancellation in (3.11). If additional operations are added, the cross terms from which the cancellation was derived here are still present. Thus Proposition 3.2 is remains valid, only with the 3 replaced by the degree of the graph. Provided this degree (= the total number of operations) is fixed, the graph still has rapid mixing.

It is unclear if including extra power operations speeds up the discrete logarithm algorithm. However, the rapid mixing of such random walks may have additional applications, such as to the stream ciphers in [8].

Acknowledgements: the authors wish to thank R. Balasubramanian, Michael Ben-Or, Noam Elkies, David Jao, László Lovász, and Prasad Tetali for helpful discussions and comments.

References

1. Eric Bach, *Toward a theory of Pollard's rho method*, Inform. and Comput. **90** (1991), 139–155.
2. Béla Bollobás, *Modern graph theory*, Graduate Texts in Mathematics, vol. 184, Springer-Verlag, New York, 1998, ISBN 0-387-98488-7.
3. J.A. Fill, *Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process*, Ann. Appl. Probab. **1** (1991), 62-87.
4. Martin Hildebrand, *A survey of results on random walks on finite groups*, Probab. Surv. **2** (2005), 33–63 (electronic).
5. Jeremy Horwitz, *Applications of Cayley Graphs, Bilinearity, and Higher-Order Residues to Cryptology*, Ph.D. Thesis, Stanford University, 2004, `http://math.scu.edu/~jhorwitz/pubs/`.
6. Jeremy Horwitz and Ramarathnam Venkatesan, *Random Cayley digraphs and the discrete logarithm*, Algorithmic Number Theory (Sydney, 2002), Lecture Notes in Comput. Sci., vol. 2369, Springer, Berlin, 2002, pp. 416–430.
7. László Lovász and Peter Winkler, *Mixing times*, Microsurveys in Discrete Probability (Princeton, NJ, 1997), DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 41, Amer. Math. Soc., Providence, RI, 1998, pp. 85–133.
8. Stephen D. Miller, Ilya Mironov, and Ramarathnam Venkatesan, *Fast and Secure Stream Cipher Designs Using Rapidly Mixing Random Walks and Revolving Buffers* (2005), Preprint.
9. Ravi Montenegro and Prasad Tetali, *Mathematical Aspects of Mixing Times in Markov Chains*, Foundations and Trends in Theoretical Computer Science, 2006.
10. Hugh L. Montgomery, *Ten lectures on the interface between analytic number theory and harmonic analysis*, CBMS Regional Conference Series in Mathematics, vol. 84, Published for the Conference Board of the Mathematical Sciences, Washington, DC, 1994, ISBN 0-8218-0737-4.
11. Victor Shoup, *Lower bounds for discrete logarithms and related problems*, Advances in Cryptology—EUROCRYPT '97 (Konstanz), Lecture Notes in Comput. Sci., vol. 1233, Springer, Berlin, 1997, pp. 256–266, Updated version at `http://www.shoup.net/papers/dlbounds1.pdf`.
12. J. Michael Steele, *The Cauchy-Schwarz master class*, MAA Problem Books Series, Mathematical Association of America, Washington, DC, 2004, ISBN 0-521-83775-8, 0-521-54677-X, An introduction to the art of mathematical inequalities.
13. Edlyn Teske, *On random walks for Pollard's rho method*, Math. Comp. **70** (2001), 809–825.
14. ______, *Square-root algorithms for the discrete logarithm problem (a survey)*, Public-Key Cryptography and Computational Number Theory (Warsaw, 2000), de Gruyter, Berlin, 2001, pp. 283–301.